arXiv:1611.10162v1 [cs.CV] 27 Nov 2016

Predicting the Category and Attributes of Mental Pictures Using Deep Gaze Pooling

Hosnieh Sattar[1,2] Andreas Bulling[1] Mario Fritz[2]
[1]Perceptual User Interfaces Group, [2]Scalable Learning and Perception Group
Max Planck Institute for Informatics, Saarbrücken, Germany
{sattar,mfritz,bulling}@mpi-inf.mpg.de

Abstract

Previous work focused on predicting visual search targets from human fixations but, in the real world, a specific target is often not known, e.g. when searching for a present for a friend. In this work we instead study the problem of predicting the mental picture, i.e. only an abstract idea instead of a specific target. This task is significantly more challenging given that mental pictures of the same target category can vary widely depending on personal biases, and given that characteristic target attributes can often not be verbalised explicitly. We instead propose to use gaze information as implicit information on users' mental picture and present a novel gaze pooling layer to seamlessly integrate semantic and localized fixation information into a deep image representation. We show that we can robustly predict both the mental picture's category as well as attributes on a novel dataset containing fixation data of 14 users searching for targets on a subset of the DeepFahion dataset. Our results have important implications for future search interfaces and suggest deep gaze pooling as a general-purpose approach for gaze-supported computer vision systems.

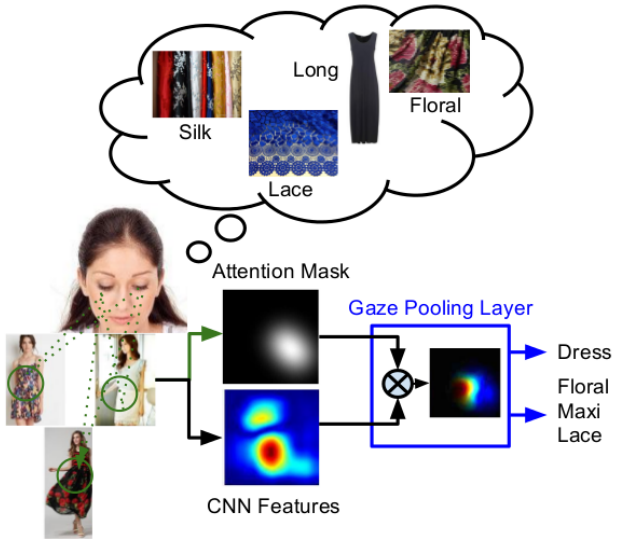

Figure 1. We propose a method to predict the mental picture – consisting of abstract ideas – from users' gaze. Our method uses gaze information as implicit information and a novel gaze pooling layer to seamlessly integrate localized fixation information on images (as attention masks) into a deep image representation.

1. Introduction

The cognitive process of mind as well as thoughts and intentions are reflected in human gaze behavior [4, 17, 19]. In his seminal work, Yarbus showed that gaze behavior is also closely related to the task that people perform [37]. The work triggered a large number of follow-up works in various fields, including artificial intelligence, neuroscience, as well as human and computer vision.

These works demonstrated that gaze behavior can be used to infer users' tasks [2, 13, 39], activities [5, 31], as well as cognitive [6] and abstract thought processes [8, 23], or it can be used for interactive image retrieval [7, 14, 27].

Most recently, gaze information was successfully used to predict the target of a visual search task [3, 12, 28, 30, 39].

Said prediction is challenging given that the target is only in the mind of the users. Borji et al. showed the possibility of search target prediction in a closed-set (all the search target were in the training set) [3] while Zelinsky et al. predicted categorical search targets from user fixations [39].

To overcome this limitation, Sattar et al. more recently introduced a method that extended search target prediction to an open-set [30]. Their method was able to predict the search target over an unknown set of targets. Al thought they did not train explicitly for all the targets but they still required a defined set of targets to measure the similarities of fixations to them.

While all of these works underline the significant potential of search target prediction using gaze, they tried to predict the specific instance of the target, which had to be

shown to users beforehand. This represents a rather strong requirement given that, in the real world, a specific target is often not known a priori, e.g. when searching for a suitable dress for a friend's upcoming birthday party. Instead, users often only have a *mental picture* in mind, i.e. only an abstract idea of what they are looking for. In addition, they might only be able to verbalise individual desired properties of a that mental picture while other properties remain unspecific. Considering the birthday example again, a user might be looking for a dress that "is flowery because I like flowers" and "it is long as I want to be classic" or "I would like it made from lace or silk".

As a proxy to the real world, similar to prior work [30], we study prediction of the mental picture during visual search in image collages. As shown in Figure 1, a user has a mental picture of what she is looking for. While she is looking at the collage, part of one image might be compatible with her mental picture while others are unrelated. As in the example, one of the dresses has the same material as observers mental picture and other has similar pattern. We then aim to predict the mental picture. This task is significantly more challenging given that mental pictures of the same target category can vary widely depending on personal biases, and given that characteristic target attributes can often not be verbalised explicitly.

The specific contributions of this work are three-fold. First, we present the first method to predict the categories and attributes of mental picture from gaze and image information. Second, we propose a novel gaze-pooling layer that seamlessly integrates fixation data with state-of-the-art deep learning architectures. We show that it is possible to add gaze information into a CNN formulation without a need to train the CNN for gaze data. Third, we present a new data set for attribute and category prediction of mental pictures, which contains fixations data of 14 users, during searching for clothing categories and attributes. We report on a series of experiments on this dataset that demonstrate both the feasibility of the task as well as the effectiveness of our method. Our results have important implications for future search interfaces and suggest deep gaze pooling as a general-purpose approach for gaze-supported computer vision systems.

2. Related Work

Our method for predicting the mental picture using gaze is related to previous works on gaze-supported computer vision, most importantly the prediction of visual search targets, as well as previous efforts to get at the mental picture. **Gaze-Supported Computer Vision.** Accessibility of the eye tracking instruments increases the interest of using gaze informations to aid computer vision tasks. Vision research shown that human tend to fixate on salient objects in the visual scene [11, 26]. Using this motivation, gaze cues were used to infer the object locations. In [20, 36] they used the correlation of fixations and object in the context of saliency predictions. In [27, 16] they used fixation data to to perform weakly supervised training of object detectors, or in [22]they inferred scene semantics. Gaze information has been used to analyze pose estimation tasks [32] as well as detecting actions [24]. Gaze has been used in segmentation as well. [25] used human fixations to do an active segmentation. Fixation information is also used for localizing important objects in egocentric videos [34, 9]. Recently, [33], studied the possibility of integrating human gaze information into an attention-based long short-term memory architecture. They showed the effectiveness of gaze in image captioning and complementing the machine attention for semantic scene understanding task.

Predicting Search Targets Using Gaze. Gaze informations has been used to predict the target of visual search task. Zelinsky et al. predicted search targets from subjects' gaze patterns during a categorical search task [39]. In their experiments, participants were asked to find two categorical search targets among four visually similar distractors. Borji et al. focused on predicting search targets from fixations [3]. They used a binary pattern and 3-level luminance patterns as target and participants were asked to find targets out of a set of other patterns. In a recent study Sattar et al. predicted search targets of users in an open and closed-world setting [30]. They introduced a method, which measured the compatibility of the gaze data and search targets. Participants were asked to find a target image in synthesized collages of images. All of these works, however, only focused on predicting a specific instance of search target fixations data.

In contrast, our work is first to address the prediction of the semantic concepts and properties of mental picture using gaze. it removes the need of seeing specific stimulus.

Getting at Mental Pictures. Visual attributes can be considered as semantic properties that users can use to address their mental picture [10, 29, 35]. Kovashka et al. introduced a novel explicit feedback method to assess the mental picture of users [1]. In their method, a user described which attributes of exemplar images should be modified via semantic language to communicate with the system. Their method updates its relevance function iteratively and re-ranks the database of images until they match the mental-picture.

Most recently, Yu et al. proposed to use free-hand human sketches as queries to perform instance-level retrieval of images [38]. They considered these sketches to be manifestations of users' mental pictures. They developed a deep triplet-ranking model for instance-level SBIR.

To overcome the insufficient fine-grained training data, they introduced a novel data augmentation and staged pre-training strategies. Both works are related to ours as they tried to predict the mental picture of users. However, [1]

required users to define their mental picture with a set of attributes and assumed that they have a large enough vocabulary to describe their mental picture. For [38], sketching may not be easy for everyone and sketching concepts like texture, color, material, style is difficult if not impossible. In contrast, in our work, we are not using any feedback loop our explicit user input and we do not require any explicit initial description of the mental picture. We, for the first time, instead use fixation information that can be acquired implicitly during the search task itself.

3. Data Collection and Collage Synthesis

To evaluate or method we used the DeepFashion dataset [21]. DeepFashion is a clothes dataset consisting of 289,222 images annotated with 46 different categories, 1000 attributes. We used the top 10 categories and attributes in our study. The training set of DeepFashion dataset was used to train the our image model for categories and attributes predictions. We used the validation set of deepfashion dataset, to train users for each category and attributes. Finally, the test set is used to build up image collages, which we record human gaze data while searching for specific categories and attributes. In the following our data collection is explained in more details.

3.1. Participants and Apparatus

We collected data from 14 participants (six females), aged between 18 and 30 years and with different nationalities. All participants had normal or corrected to normal vision. For gaze data collection we used a stationary Tobii TX300 eye tracker that provides binocular gaze data at a sampling frequency of 300Hz. Parameters for fixation detection were left at their defaults: fixation duration was set to 60ms while the maximum time between fixations was set to 75ms. The stimuli were shown on a 30-inch screen with a resolution of 2560x1600 pixels. Before the recording, we calibrated the eye tracker using a standard 9-point calibration, followed by a validation of eye tracker accuracy. Participants were assigned to search for ten different categories and attributes in two different sessions: training and testing.

3.2. User Training

We first trained participants so they could familiarized themselves with the categories and attributes. To this end, for each category, each participant was first shown nine randomly selected samples of that category. They were then shown four randomly selected samples, two of which belonged to the category they were currently being trained for and the other two from the remaining categories. For each sample, participants had to answer whether it belonged to the current category or not. After entering their response, participants received visual feedback on the right answer.

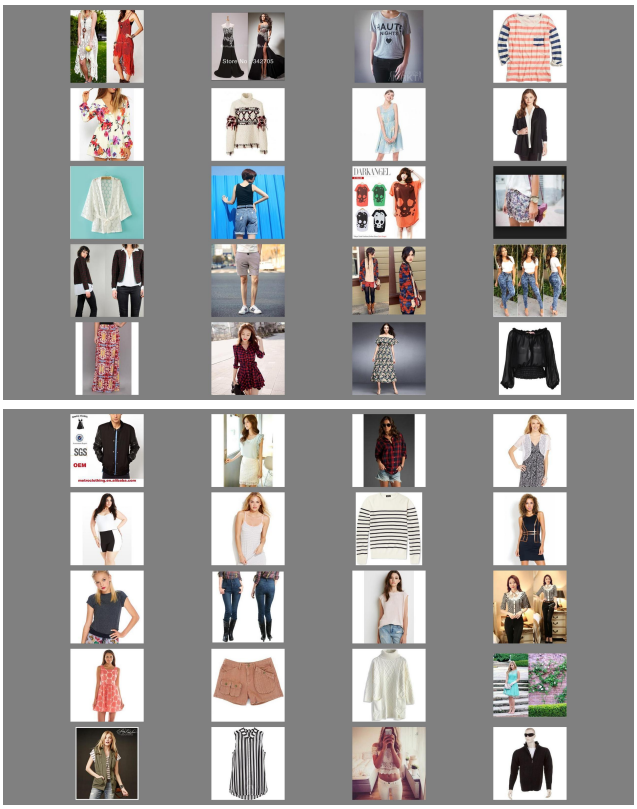

Figure 2. Sample image collages used for data collection: Attributes (top), Categories (bottom). Participants were asked to find different clothing attributes and categories within these collages.

Participants had to provide at least five correct answers before proceeding to the next category or attribute. In addition to these binary answers, participants had to complete several multiple-choice questions. Participants were shown five samples from each category or attributes and had to select the right category. This step is repeated two times for all of the categories and attributes.

3.3. User Testing

Afterwards, we recorded participants' gaze while they searched for each of the categories or attributes in image collages. For each category and attribute we generated 10 image collages, each containing 20 images. Each category or attribute appeared only twice in each collage at a random location (see Figure 2 for an example). Collages were shown full screen. Participants were then shown three different samples of the first out of 10 categories or attributes. Participants had a maximum of 10 seconds to find the target category or attribute in the collage. To determine more easily on which images participants fixated on, all images were placed on a grey background and had a margin to neighboring images of at least 60 pixels. As soon as participants found the target they were asked to press a key. Afterwards they were

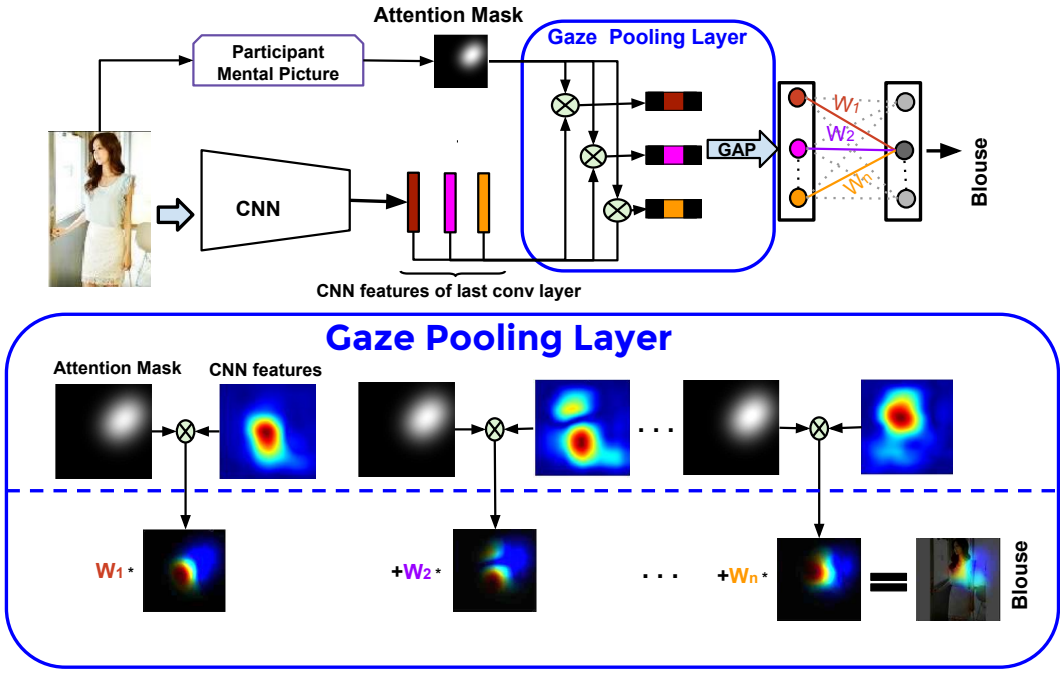

Figure 3. (top) Overview of our approach to predicting attributes and categories of mental pictures. The input image is encoded via a CNN architecture into multiple spatial feature layers. A participant fixates on the same image with a mental picture in mind. Our Gaze Pooling Layer combines visual featuers and gaze attention mask in a feature weighting scheme. The output is classified into categories or attributes of the mental picture.(bottom) Illustration of the proposed Gaze pooling layer that combines attention masks with CNN feature maps via a spatial reweighting.

asked whether they had found the target and how difficult the search had been. This procedure was repeated ten times for ten different categories or attributes, resulting in a total of 100 search tasks.

4. Prediction of Mental Pictures Using Gaze

In this work we are interested in predicting the category and attributes of mental pictures. Previous works either required users' explicit input to get their mental picture [38, 1] or only predicted a specific search target instance shown to the users beforehand [3, 30, 39]. In contrast, we introduce a novel approach that uses gaze input as implicit information on the mental picture. We address this task by introducing the Gaze Pooling Layer (GPL) In the following, we describe the four major components of our method in detail: The image encoder, human gaze encoding, the Gaze Pooling Layer, and Classification. Finally, we also discuss different integration schemes across multiple images.

4.1. Image Encoder

We build on the recent success of deep learning and use a convolutional neural network (CNN) to encode image information [15, 18]. Given a raw image I, a CNN is used to extract image feature map $F(I)$.

$$F(I) = \textbf{CNN}(I)$$

Due to the end-to-end training properties of these networks, this allows us to obtain domain specific features. In our case, the network will be trained with data and labels relevant to the fashion domain.

As we are interested in combining spatial gaze features with the image features, we use features $F(I)$ of the last convolutional layer that still has a spatial resolution. This results in a task-dependent representation with spatial resolution. In addition, to gain a higher spatial resolution we used a similar architecture as describe in[41]. For AlexNet the last layers after `conv5` are removed. The resulting resolution is 13×13. For VGGnet, layers after `conv5-3` are removed to gain a mapping resolution of 14×14.

4.2. Human Gaze Encoding

Given a mental picture, participant $P \in \mathbb{P}$ look at image I and performs fixations $G(I,P) = (x_i, y_i), i = 1, ..., N$ in screen coordinates. We aggregate these fixations into attention masks that capture the spatial density $A(G)$ of fixations over the full image. Therefore, we represent the attention mask $A(g)$ for a single fixations $g \in G(I, P)$ by a Gaussian:

$$A(g) = \mathcal{N}(g, \sigma_{\text{fix}})$$

, centered at the coordinates of the fixation, with a fixed standard deviation σ_{fix} – the only parameter of our representation. The attention mask for all fixations $A(G)$ is obtained by coordinate-wise summation:

$$A(G) = \sum_{g \in G} A(g)$$

This corresponds to an average pooling integration. We also propose a max pooling version as follows:

$$A(G) = \max_{g \in G} A(g)$$

4.3. Gaze Pooling Layer

We combine the visual features $F(I)$ with attention mask $A(G)$ in a Gaze Pooling Layer. The integration is performed by element-wise multiplication between both to obtain a gaze-weighted feature map (GWFM)

$$\mathit{GWFM}(I, G) = F(I) \otimes A(G).$$

In spirit of [41], we then perform Global Average Pooling (GAP) on each feature channel separately in order to yield a vector-valued feature representation.

$$\mathit{GAP}_{\text{GWFM}}(I, G) = \sum_{x,y} \mathit{GWFM}(I, G)$$

We finish our pipeline by classification with a fully connected layer and a softmax layer.

$$p(C|I, G) = \mathit{softmax}(W\ \mathit{GAP}_{\text{GWFM}}(I, G) + b),$$

where W are the learned weights and b is the bias. The classes represent either categories or attributes depending on the experiment and we decide for the class with the highest class posterior.

4.4. Integration across Images

In our study, a stimulus is a collage with a set of images $I_i \in \mathbb{I}$. During the search task, participants are free to fixate on multiple images in the collage, which generates fixations $G_i \in \mathbb{G}$ for each image I_i. Hence, we need a mechanism to aggregate information across images. To do this, we propose a simple but effective weighted average scheme of the computed posteriors:

$$p(C|\mathbb{I}, \mathbb{G}) = \sum_i \frac{w_j}{\sum_j w_j} p(C|I_i, G_i).$$

We consider for the weights w_i the total fixations duration on image I_i as well as fixed w_i, which leads, again, to plain averaging.

Vision-Only	——— Category ———			Attributes
Models	Top1	Top3	Top5	Accuracy
Alexnet-GAP	62%	86%	94%	75%
VGGnet-GAP	**68%**	**89%**	**96%**	**79%**

Table 1. Evaluation of vision-only methods.

4.5. Vision-Only Performance

To obtain vision-only models with higher spatial feature map resolution, we followed [41] and built on the Alexnet-GAP and VGGnet-GAP models. For our categorization experiments we fine-tuned on a 10 class classification problem on the DeepFashion dataset [21]. For attribute prediction, we fine-tuned a model with 10 times 2-way classification in the final layer. In Table 1, we report the obtained accuracies. As the VGG-based version gives superior performance across all metrics in both tasks, we used it in all further experiments.

We also performed a validation of our model in the same setting as [21] and obtained comparable results ($\pm 5\%$). Please note that all of these results are on a 10 class subset relevant to our main experiments. It also has to be noted that those numbers are not comparable to the results in the main experimental section as they represent recognition of the image content and are *not* related to the mental picture of the observer – which is the task in the following section. In fact, we conducted sanity checks that the vision-only baselines performed at chance level on all following experiments on category and attribute prediction of mental pictures.

5. Experiments

To evaluate our method for mental picture category and attribute prediction, we performed a series of experiments on the recorded dataset. We first evaluated the effectiveness of our Gaze-Pooling Layer, the effect of using a local vs a global representation, and of using a weighting by fixation duration. We then evaluated the gaze encoding that encompasses the pooling scheme of the individual fixation as well as the σ_{fix} parameter to represent a single fixation. Finally, we evaluated the robustness of our method with respect to noise in the eye tracking data, which sheds light on different possible deployment scenarios and hardware that our approach is amendable to. Across the results, we present Top-N accuracies denoting correct predictions if the correct answer is among the top N predictions.

5.1. Evalution of Gaze-Pooling Layer (GPL)

Fixation information enters our GPL in two places: The attention masks and the weighted average across images in the form of fixation duration. We first evaluated the effect of both on overall performance. We refer to the approach that

Global vs. Local	◔	Category Top1	Category Top2	Category Top3	Attribute Accuracy
Global		31%±5	48% ±8	62% ±8	59%±2
Local		49%±7	68%±6	78%±6	72 %±2
Global	✓	52%±6	68%±6	78%±6	73%±3
Local	✓	**57%±8**	**74%±7**	**84%±4**	**77%±9**

Table 2. Evaluation of global vs. local gaze pooling with and without weighting based on the fixation duration ◔.

uses the attention masks as *local* (using spatial fixation information), while the *global* approach instead uses a uniform and therefore uninformative attention mask. In addition to weighting by fixation duration we considered treating each fixation equally.

Table 2 shows the result of this evaluation, with the first column denoting if local or global information was used and the second column ◔ whether fixation duration was used. Absolute performance of our best model using local information and fixation duration were 57%, 74%, and 84% on top1-3 accuracy respectively for the categorization task and 77% accuracy for attributes. The results show a consistent improvement (16 to 18 pp for categories, 13 pp for attributes) across all measures and tasks going from a global to a local representation (first to second row). Adding the weighting by fixation duration yields another consistent improvement for both local and global approach (another 6 to 21 pp for categories). Our best method, improves overall by 22 to 26 pp on the categorization task and 18 pp on the attributes. The global method without fixation duration (first row) is in spirit similar to [30] – although the specific application differs. All further experiments will consider our best model (last row) as the reference and justify the parameter choices (average pooling, $sigma_{\text{fix}}$) by varying each parameters one by one.

5.2. Evaluation of Gaze Encoding

We then evaluated the gaze encoding that takes individual fixations as input and produces an attention mask. We first evaluated the representation of a single fixation, which depends on the parameter σ_{fix}, and then the pooling scheme that combines multiple fixations into the attention mask.

Effects of Fixation Representation Parameter f_σ. The parameter σ_{fix} controls the spatial extend of a single fixation in the attention masks as described in subsection 4.2. We determined an appropriate setting of this parameter to be $\sigma_{\text{fix}} = 1.6$ in a pilot study to roughly match the eye tracker accuracy and analyzed here the influence on the overall performance by varying this parameter in a sensible range (given eye tracker accuracy and coarseness of feature map) from 1 to 2 as shown in Table 3. As can be seen from the Table, our method is largely insensitive to the investigated range of reasonable choices of this parameter and our choice of 1.6 is on average a valid choice within that range.

$\sigma_{\text{fix}} \rightarrow$	1	1.2	1.4	1.6	1.8	2
Top1	55%	54%	56%	56%	57%	57%
Top2	74%	74%	74%	74%	74%	75%
Top3	83%	84%	84%	85%	85%	84%

Table 3. Evaluation of different gaze encoding schemes using different per-fixation σ_{fix}.

Fixation Pooling	Category Top1	Category Top2	Category Top3	Attribute Accuracy
Max	54%±8	73%±9	83%±6	76%±11
Average	**57%±8**	**74%±7**	**84%±4**	**77%±9**

Table 4. Evaluation of different fixation pooling strategies using average or max pooling.

Fixation Pooling Strategies We evaluated two options for how to integrate single fixations into an attention mask: Either using average or max pooling. The results are shown in Table 4. As the Table shows, while both options perform well, average pooling consistently improves over the max pooling option.

5.3. Noise Analysis

While our gaze data is recorded with a highly-accurate stationary eye tracker, there are different modalities and types of eye trackers available. One key characteristic in which they differ is the error at which they can record gaze data – typically measured in degrees of visual angle. While our controlled setup provides us with an accuracy of about 0.7 degrees of error, state-of-the-art eye trackers based on webcams, tablets or integrated into glasses can have up to 4 degrees depending also on the deployment scenario [40]. Therefore, we finally investigated the robustness of our approach w.r.t. different levels of (simulated) noise in the eye tracker. To this end, we sampled noise from a normal distribution with $\sigma = 1, 3, 5$. This corresponds roughly to $60, 120$ and 200 pixels and to $1.2, 2.5$ and 4.2 degrees of visual angles and hence covers a realistic range of errors.

The results of this evaluation are shown in Figure 4. As can be seen, our method is quite robust to noise with only a

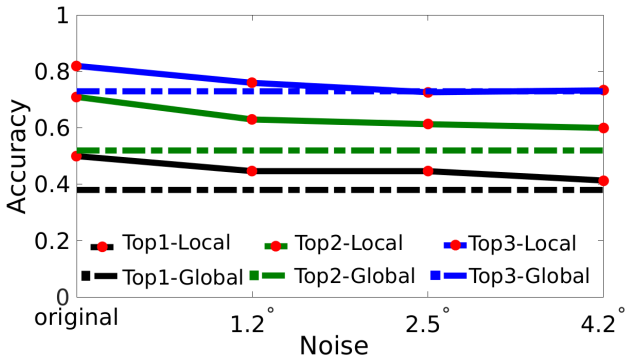

Figure 4. Accuracy for different amounts of noise added to the eye tracking data. Our method is robust to this error which suggests that it can also be used with head-mounted eye trackers or learning-based methods that leverage RGB cameras integrated into phones, laptops, or public displays.

drop of 5 to 10pp for Top3 to Top1 accuracy, respectively – even at the highest noise level. In particular, all the results are consistently above the performance of the corresponding global methods shown as dashed lines in the plot.

6. Disscussion

In this work we studied the problem of predicting the mental picture during visual search task from fixations. Table 2 shows that we can predict the mental picture significantly above chance for categories ($Top3 = 84\%$) and attributes (77%). Our Gaze Pooling Layer achieves a modular integration of visual and gaze features that is compatible with modern deep learning architectures. Beyond this capability, we would like to highlight three features that are of particular practical importance.

Parameter Free Integration Scheme First, our proposed integration scheme is parameter-free. We introduce a single parameter σ_{fix} but the gaze encoding is only input to the integration scheme and, in addition, the method turns out to be not very sensitive to the choice (see experiments in subsection 4.2).

Training from Vision-Only Input Second, fixing the attention masks to uniform maps yields a deep architecture similar to a GAP network that is well-suited for various classification tasks. While this no longer addresses the task of predicting categories and attributes intended by the human in the loop, it allows us to train the remaining architecture for the task at hand and on vision-only data, which is typical easier to obtain in larger quantities than gaze data. This type of training results in a domain-specific image encoding as well as task-specific classifier.

Training Free Gaze Deployment Gaze data is time consuming to acquire – which makes it rather incompatible with today's data hungry deep learning models. However, in our model, the attention masks computed from the gaze data can be understood as spatially localized feature importance that are used to weight feature importance in the spatial image feature maps. It turns out that this re-weighting scheme does not require re-training with gaze data (see experiments), so that the approach can be instantly deployed based on the vision-only model without any gaze-specific training. We believe this is a key feature to that makes the use of gaze information in deep learning practical.

Visualizing Effect of Gaze Pooling Layer We provide further insights into the working of our Gaze Pooling Layer by showing visual examples of the classification process and associated attention mask. Figure 5 shows results for the categorization task and Figure 6 for the attribute task.

Each of these figures shows the output of the "vision only" method that does not use gaze information and is therefore "unaware" of the mental picture of the participant. Consequently, all predictions are wrong, as they do not related to the mental picture. Instead, we observe that in each case a prominent part of the image is correctly classified as the contained category or attribute.

Next to the "vision only" prediction, we show the prediction using gaze information via our proposed Gaze Pooling Layer. We observe that in all cases we manage to predict the correct category of the mental picture and in 3 out of 4 cases at least one relevant attribute. Key to our success is the attention masks visualized in the left half of the images. It shows the selectiveness towards more relevant parts in the image that in turn re-weight visual features in our Gaze Pooling Layer.

In order to illustrate the challenges our Gaze Pooling Layer has to deal with in terms of the variations in the observed gaze data, we show example fixation data in Figure 7. In each image, fixation data of two participants (red and green dots) is overlaid over a presented collage. Although both participants had the same mental picture (top: attribute 'Floral'; bottom: category 'Cardigan'), we observe a drastically different fixation behaviour. Despite this strong variation in the gaze information, our Gaze Pooling Layer allows to predict the correct answer in all 4 cases. The key to this success is aggregating relevant local visual feature across all images in the collage, that in turn represent one consistent mental picture in terms of categories and attributes.

7. Conclusion

In this work we demonstrated how to predict the mental picture of users from their fixations data during visual search. To this end we proposed a novel gaze pooling layer to seamlessly integrate semantic and localized fixation information

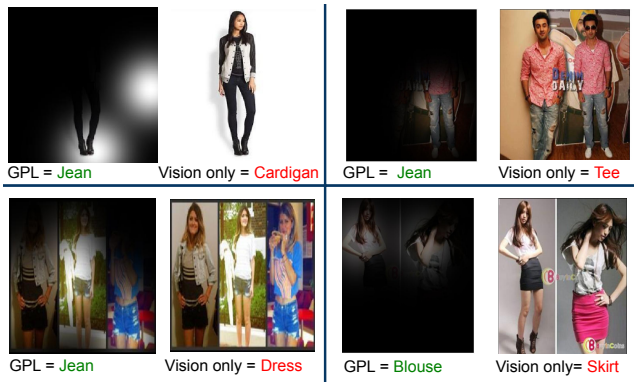

Figure 5. Example category responses of Gaze Pooling Layer(GPL) and Vision only method.

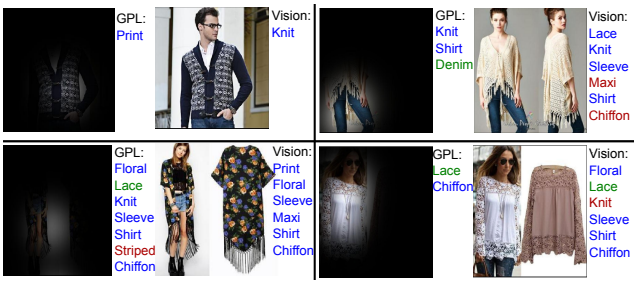
GPL: Print
Vision: Knit
GPL: Knit Shirt Denim
Vision: Lace Knit Sleeve Maxi Shirt Chiffon
GPL: Floral Lace Knit Sleeve Shirt Striped Chiffon
Vision: Print Floral Sleeve Maxi Shirt Chiffon
GPL: Lace Chiffon
Vision: Floral Lace Knit Sleeve Shirt Chiffon

Figure 6. Example attribute responses of Gaze Pooling Layer(GPL) and Vision only method.

into a deep image representation. Our work is fundamentally different to previous ones as we no longer require participants to find specific instance of an object. Also, we are not training on gaze data which makes our method generic and generalizable to deployments in which gaze is not available. Our work paves the way for new applications in different areas, including gaze-assisted computer vision, visual question answering, attention-based convolutional neural networks, and context-based image retrieval.

References

[1] K. G. A. Kovashka, D. Parikh. Whittlesearch: Interactive image search with relative attribute feedback. *International Journal of Computer Vision*, 115(2), 2015. 2, 4

[2] A. Borji and L. Itti. Defending yarbus: Eye movements reveal observer's task. *Journal of vision*, 14(3), 2014. 1

[3] A. Borji, A. Lennartz, and M. Pomplun. What do eyes reveal about the mind?: Algorithmic inference of search targets from fixations. *Neurocomputing*, 2014. 1, 2, 4

[4] F. J. Brigham, E. Zaimi, J. J. Matkins, J. Shields, J. McDonnouugh, and J. J. Jakubecy. *The eyes may have it:Reconsidering eye-movement research in human cognition*, volume 15. Advances in Learning and Behavioral Disabilities, 2001. 1

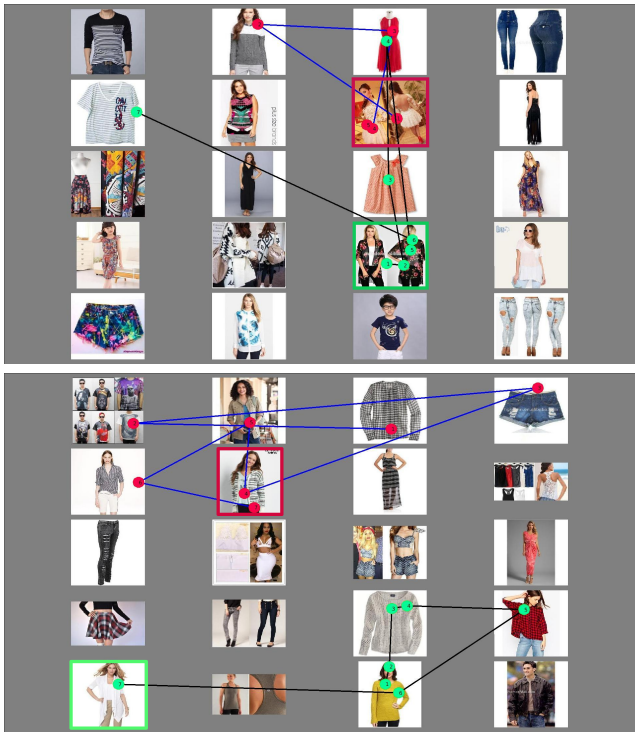

Figure 7. Example fixation data of 2 participants (red and green dots) with mental picture attribute='Floral' on top and category='Cardigan' below.

[5] A. Bulling, J. A. Ward, H. Gellersen, and G. Tröster. Eye movement analysis for activity recognition using electrooculography. *IEEE TPAMI*, 33(4), 2011. 1

[6] A. Bulling and T. O. Zander. Cognition-aware computing. *IEEE Pervasive Computing*, 13(3), 2014. 1

[7] J. Coddington, J. Xu, S. Sridharan, M. Rege, and R. Bailey. Gaze-based image retrieval system using dual eye-trackers. In *ESPA*, 2012. 1

[8] R. Coen-Cagli, P. Coraggio, P. Napoletano, O. Schwartz, M. Ferraro, and G. Boccignone. Visuomotor characterization of eye movements in a drawing task. *Vision Research*, 49(8), 2009. 1

[9] D. Damen, T. Leelasawassuk, O. Haines, A. Calway, and W. Mayol-Cuevas. You-do, i-learn: Discovering task relevant objects and their modes of interaction from multi-user egocentric video. In *BMVC*, 2014. 2

[10] E. E. Zavesky and S. Chang. Cuzero: Embracing the frontier of interactive visual search for informed users. *ACM-MM*, 2008. 2

[11] W. Einhäuser, M. Spain, and P. Perona. Objects predict fixations better than early saliency. *Journal of Vision*, 8(14), 2008. 2

[12] A. Haji-Abolhassani and J. J. Clark. A computational model for task inference in visual search. *Journal of Vision*, 13(3), 2013. 1

[13] A. Haji-Abolhassani and J. J. Clark. An inverse yarbus process: Predicting observer's task from eye movement patterns. *Vision research*, 2014. 1

[14] Z. Hussain, A. Klami, J. Kujala, A. P. Leung, K. Pasupa, P. Auer, S. Kaski, J. Laaksonen, and J. Shawe-Taylor. Pinview: Implicit feedback in content-based image retrieval. arXiv 1410.0471, 2014. 1

[15] A. V. K. Simonyan, Karen and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv:1312.6034*, 2013. 4

[16] S. Karthikeyan, V. Jagadeesh, R. Shenoy, M. Ecksteinz, and B. Manjunath. From where and how to what we see. In *ICCV*, 2013. 2

[17] C. L. Kleinke. Gaze and Eye Contact: A Research Review. *Psychological Bulletin*, 100(1), 1986. 1

[18] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 4

[19] M. F. Land and S. Furneaux. The knowledge base of the oculomotor system. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 352(1358), 1997. 1

[20] Y. Li, X. Hou, C. Koch, J. M. Rehg, and A. L. Yuille. The secrets of salient object segmentation. In *CVPR*, 2015. 2

[21] Z. Liu, P. Luo, S. Qiu, X. Wang, and X. Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *CVPR*, 2016. 3, 5

[22] E. Marinoiu, D. Papava, and C. Sminchisescu. Pictorial human spaces. how well do humans perceive a 3d articulated pose? In *ICCV*, 2013. 2

[23] F. W. Mast and S. M. Kosslyn. Eye movements during visual mental imagery. *Trends in Cognitive Sciences*, 6(7), 2002. 1

[24] S. Mathe and C. Sminchisescu. Multiple instance reinforcement learning for efficient weakly-supervised detection in images. *arXiv:1412.0100*, 2014. 2

[25] A. Mishra, Y. Aloimonos, and C. L. Fah. Active segmentation with fixation. In *ICCV*, 2009. 2

[26] A. Nuthmann and J. M. Henderson. Object-based attentional selection in scene viewing. *Journal of vision*, 10(8), 2010. 2

[27] G. Papadopoulos, K. Apostolakis, and P. Daras. Gaze-based relevance feedback for realizing region-based image retrieval. *ACM-MM*, 16(2), Feb 2014. 1, 2

[28] U. Rajashekar, A. C. Bovik, and L. K. Cormack. Visual search in noise: Revealing the influence of structural cues by gaze-contingent classification image analysis. *Journal of Vision*, 6(4), 2006. 1

[29] M. Rastegari, A. Diba, D. Parikh, and A. Farhadi. Multi-attribute queries: To merge or not to merge? In *CVPR*, 2013. 2

[30] H. Sattar, S. Müller, M. Fritz, and A. Bulling. Prediction of search targets from fixations in open-world settings. In *CVPR*, 2015. 1, 2, 4, 6

[31] J. Steil and A. Bulling. Discovery of everyday human activities from long-term visual behaviour using topic models. In *UbiComp*, 2015. 1

[32] R. Subramanian, V. Yanulevskaya, and N. Sebe. Can computers learn from humans to see better?: inferring scene semantics from viewers' eye movements. In *ACM-MM*, 2011. 2

[33] Y. Sugano and A. Bulling. Seeing with humans: Gaze-assisted neural image captioning. arxiv:1608.05203, 2016. 2

[34] T. Toyama, T. Kieninger, F. Shafait, and A. Dengel. Gaze guided object recognition using a head-mounted eye tracker. In *ETRA*, 2012. 2

[35] P. N. B. W. Scheirer, N. Kumar and T. E. Boult. Multi-attribute spaces: Calibration for attribute fusion and similarity search. In *CVPR*, 2012. 2

[36] J. Xu, M. Jiang, S. Wang, M. S. Kankanhalli, and Q. Zhao. Predicting human gaze beyond pixels. *Journal of vision*, 14(1), 2014. 2

[37] A. L. Yarbus. *Eye movements and vision*. Springer, 1967. 1

[38] Q. Yu, F. Liu, Y.-Z. Song, T. Xiang, T. M.Hospedales, and C. C. Loy. Sketch me that shoe. In *CVPR*, 2016. 2, 3, 4

[39] G. J. Zelinsky, Y. Peng, and D. Samaras. Eye can read your mind: Decoding gaze fixations to reveal categorical search targets. *Journal of Vision*, 13(14), 2013. 1, 2, 4

[40] X. Zhang, Y. Sugano, M. Fritz, and A. Bulling. Appearance-based gaze estimation in the wild. In *CVPR*, 2015. 6

[41] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Learning deep features for discriminative localization. In *CVPR*, 2016. 4, 5